

SAMGRID

on the General Purpose Farm

Lessons Learned and prospects for wider
implementation.

Steven Timm
for CSS/SCS
GDM Aug 31 2004

History of SAMGRID Testbed

- Established in March of 2004
 - Purpose was for SAM-GRID and JIM Developers to test out software
 - D0 Runtime Environment tests also requested but haven't happened to our knowledge
 - These nodes are part of General Purpose farms
 - This was requested so that development could occur in a “generic” environment with nothing D0 specific
 - Development has centered on D0 MC and CDF MC thus far.
-
-

SAMGRID Test bed setup

- 31 worker nodes, formerly from the old D0 farm
 - 13 more nodes still coming
 - 1 grid interface node “samgfarm2” runs Globus gatekeeper
 - The other grid interface node “samgfarm1” is SAM station, supplying data to and from the worker nodes.
 - Local batch manager is FBSNG, hosted on fnsfo
 - Testbed nodes are reserved for SAMGRID activities only.
Only SAMGRID user can run on them, and SAMGRID user can't run anywhere else in the farm.
-
-

Authentication and Authorization

- In on-site FBSNG operation we require Kerberos credential to submit a job to FBSNG.
 - In grid operation, incoming jobs from the grid are mapped to userid's sam or samgrid.
 - Users sam and samgrid are allowed to submit jobs to FBSNG only from trusted head nodes, without kerberos credentials.
 - Only KCA certificates accepted at the moment. To accept others we would need strong auth. Waiver.
 - Only SAMGRID developers running jobs at the moment.
 - In initial discussion, it was proposed to add x.509 authentication to FBSNG. This has not been done as yet.
-
-

Demand

- Main General purpose farm users:
 - SDSS, KTEV, MINOS, MiniBOONE, Auger, E871, Accelerator division, EPP Theory, EPP Astro
 - To date none of them have asked for Grid access
 - We could open up current samgrid setup to these other users but probably will, instead, make a different grid head node for production use, leaving the testbed so developers can bring it up and down at will.
 - Current testbed setup is also being cloned on the D0 farms at the moment. (See Diesburg talk).
-
-

Throughput Requirements

- Current testbed design requires all cached data to go through single-node SAM station.
 - Even at 30-node level test, testbed head node samgfarm2 is often overloaded.
 - We need to get real numbers for anticipated data flow for SAM station/Globus gatekeeper.
 - Need to rightsize the machines that will be the production versions of samgfarm1 and samgfarm2
 - Anticipating to replace IRIX head nodes on General Purpose Farms with Linux head nodes next year.
-
-

Data Path Architecture

- All tape drives were recently disabled on GP Farms
 - Canonical way is to use encp on head nodes to read and write data from and to Enstore
 - dcap/dccp is available on worker nodes but lightly used...is the stken dcache up to it?
 - We need to study—should we get users to use a native dccp-like data path or get them to transition to a Grid-like path?
-
-

Administration

- In Fermi-Grid how often will grid map files be changed
 - How often will new uid's be created
 - What if any local disk space will be needed?
 - x.509 authentication in FBSNG would be nice but would add overhead of putting globus software, grid mapfiles, etc. on all nodes.
 - Current tools we have to distribute these are rsync-based package currently used on the farms, and 411 package that comes with NPACI-Rocks.
 - These tools should be adequate but we need to understand the frequency with which they will be used.
-
-

OS Version

- GP Farms scheduled to upgrade to Fermi Linux LTS 3.0.x by the end of the calendar year.
- Will the various grid software be ready?

Network Topology

- Open vs. private
 - Kerberos, dcache, encp, fbsng, dfarm, and all our management tools expect nodes on the open network
 - Various grid-related software can deal with and even prefers private networks.
 - Current farms, GP, D0, CDF, all have identical OS installation but can't interoperate due to network topology of the Fermi net.
 - Now is time to plan to make sure this doesn't happen in grid era.
-
-

Training

- We have requested a “Grid for Sysadmins” seminar, hopefully sometime this fall.
- Goal is to help work through the jargon and let system administrators really see how their systems are impacted by opening to the Grid.

